arXiv:1610.05825v1 [cs.CC] 18 Oct 2016

NP-hard sets are not sparse unless P = NP: An exposition of a simple proof of Mahaney's Theorem, with applications

Joshua A. Grochow*

October 20, 2016

Abstract

Mahaney's Theorem states that, assuming P ≠ NP, no NP-hard set can have a polynomially bounded number of yes-instances at each input length. We give an exposition of a very simple unpublished proof of Manindra Agrawal whose ideas appear in Agrawal–Arvind ("Geometric sets of low information content," Theoret. Comp. Sci., 1996). This proof is so simple that it can easily be taught to undergraduates or a general graduate CS audience—not just theorists!—in about 10 minutes, which the author has done successfully several times. We also include applications of Mahaney's Theorem to fundamental questions that bright undergraduates would ask which could be used to fill the remaining hour of a lecture, as well as an application (due to Ikenmeyer, Mulmuley, and Walter, arXiv:1507.02955) to the representation theory of the symmetric group and the Geometric Complexity Theory Program. To this author, the fact that sparsity results on NP-complete sets have an application to classical questions in representation theory says that they are not only a gem of classical theoretical computer science, but indeed a gem of mathematics.

1 Introduction

Mahaney's Theorem [Mah82] is one of the seminal results in the pre-probabilistic era of computational complexity, and answers several foundational questions about the nature of the P versus NP question. The theorem states that there are no sparse NP-complete sets unless P = NP. Recall that a set $L \subseteq \Sigma^*$ is (polynomially) sparse if the number of strings of length n in L is bounded by a polynomial in n.

Unfortunately, this theorem is much less well-known than it should be, partially due to the fact that Mahaney's original proof—using the so-called "census technique" originated by Fortune [For79]—although not complicated, takes at least one full lecture to teach. Given the myriad topics—most of them more modern and in vogue—which teachers feel they must cover in undergraduate and graduate courses on complexity, it's unclear if the census technique really deserves a whole lecture (or more), especially since the applications of the census technique more or less stop at Mahaney's Theorem. On looking at the topics often taught in complexity classes nowadays, the general consensus seems to be that it is not worth a whole lecture in lieu of other topics.

In this note, we give an exposition of a much simpler proof, due to Manindra Agrawal (unpublished, but the idea of the proof is present in Agrawal and Arvind [AA96], where the technique was put towards proving a stronger statement). This proof is so simple that it can be taught to undergraduates or a general graduate CS audience—not just theorists!—in about 10 minutes, which the author has done successfully several times. (Not to mention the value of having multiple and simpler proofs of important theorems.)

*Department of Computer Science, University of Colorado at Boulder, joshua.grochow@colorado.edu and Santa Fe Institute, jgrochow@santafe.edu

The rest of a lecture could then be spent on other topics, or, I suggest, on applications of Mahaney's Theorem to fundamental questions about P versus NP that a bright undergraduate would ask (which I cover in Section 3). Alternatively, these applications could potentially be given as homework after having taken the 10 minutes to teach Mahaney's Theorem. We'll cover more of the history in Section 3 in the course of discussing the applications of Mahaney's Theorem, but for now let's jump right into the proof. We also note that slight variants of Agrawal's proof can be used to give similarly simple proof of Fortune's Theorem—no coNP-complete set is sparse unless $\mathsf{P} = \mathsf{NP}$ (see Remark 2.2)—and a common generalization of Mahaney's and Fortune's Theorems to a slightly more general kind of reduction (see Question 3.2). The proof is reminiscent of the "left set" technique of Ogiwara–Watanabe [OW91] (which they used to generalize Mahaney's Theorem to $\leq^p_{btt}$ reductions: non-adaptive Turing reductions that make only $O(1)$ queries), but when applied to the case of Mahaney's Theorem, Agrawal's proof is the simplest proof of which the author is aware (at the very least, requiring the fewest additional definitions and lemmas).

Acknowledgments

We gratefully acknowledge Manindra Agrawal for allowing us to publish this proof. Certainly I am not the only person aside from Agrawal to have been aware of this proof; at the very least Thomas Thierauf was also aware, and was very supportive of writing it up. We thank Thierauf for his help, support, and suggestions. We also thank Lance Fortnow for useful comments on a draft.

2 The proof

Theorem 2.1 (Mahaney [Mah82])**.** *If* $\mathsf{P} \neq \mathsf{NP}$*, then no* NP*-hard set is sparse.*

Proof (M. Agrawal, compare Agrawal–Arvind [AA96]). Suppose L is a sparse NP-hard language. We will show how to solve SAT in polynomial time. Suppose we wish to decide whether a given propositional formula $\varphi(x_1, \ldots, x_n)$ is satisfiable. Consider the downward self-reduction tree of φ: in its first layer, we get the two formulas $\varphi_0 = \varphi(0, x_2, \ldots, x_n)$ and $\varphi_1 = \varphi(1, x_2, \ldots, x_n)$. Each level ℓ of the tree has the crucial property that

$$\varphi \text{ is satisfiable } \Leftrightarrow \text{ at least one formula at level } \ell \text{ is satisfiable.} \tag{1}$$

The problem with using the downward self-reduction tree to solve SAT is that it takes exponential time, since it essentially amounts to trying all 2^n possible assignments. However, if we could prune the tree at each level in polynomial time in such a way as to preserve property (1), but only have $\text{poly}(n)$ many strings at each level, then the resulting SAT algorithm would work in polynomial time, since there are only n levels. This is exactly what we will show how to do.

Let $f\colon \Sigma^* \to \Sigma^*$ be a polynomial-time many-one reduction from SAT to L. Let $s(n)$ ("s" for "sparsity") be a polynomial bounding the number of strings in L of length $\leq n$, and let $r(n)$ ("r" for "reduction") be a polynomial bound on the length-stretch of f, that is, $|f(\varphi)| \leq r(|\varphi|)$ for all φ. We will prune the downward self-reduction tree to have at most $t(n)+1$ formulas at each level ("t" for "tree"), where $t(n) = s(r(2n+5))$, which is polynomially bounded. (We use $2n + 5$ because we will apply f to strings of the form $(\varphi) \vee (\psi)$ where $|\varphi|, |\psi| \leq n$, and we had to add the 5 symbols "$() \vee ()$" to combine these two formulae together. For clarity below, we will not include these extra parentheses.)

The construction proceeds in stages, starting with the first stage at the first level of the tree. Let $\varphi_1, \ldots, \varphi_k$ be the formulas at the current level of the tree at the current stage. If $k \leq t(n) + 1$ we continue to the next level; otherwise, $k > t(n) + 1$. For $i = 2, \ldots, k$, let $q_i = f(\varphi_1 \vee \varphi_i)$, where as usual "$\vee$" denotes "OR;" note that we do not include $i = 1$. After each stage, the remaining φ_i at the current level of the tree are re-indexed to start from 1 again.

Case 1: All q_i are distinct strings. Remove φ_1 from the tree and continue to the next stage. In this case, φ_1 must have been unsatisfiable, so removing it is okay, i. e., preserves property (1): since there are strictly more than $s(r(2n+5))$ distinct q_i, at least one of them must not be in $L \cap \Sigma^{\leq r(2n+5)}$, say q_i. Since $q_i = f(\varphi_1 \vee \varphi_i) \notin L$, $\varphi_1 \vee \varphi_i$ is not satisfiable, and in particular φ_1 is not satisfiable. (Note that we do not know, and we do not *need* to know, which q_i is not in L.)

Case 2: $q_i = q_j$ for some distinct i, j. Remove φ_i from the tree and continue to the next stage. Let's see why this is okay. If φ_1 is satisfiable, then removing φ_i preserves property (1) (recall $i \neq 1$). If φ_1 is not satisfiable, then φ_i is satisfiable if and only if $\varphi_1 \vee \varphi_i$ is satisfiable, if and only if $\varphi_1 \vee \varphi_j$ is satisfiable (since $f(\varphi_1 \vee \varphi_i) = f(\varphi_1 \vee \varphi_j)$), if and only if φ_j is satisfiable. So again, removing φ_i preserves property (1). □

Remark 2.2. One can give a similarly simple proof of Fortune's Theorem [For79], which says that no coNP-hard set is sparse unless P = NP, or equivalently that no NP-hard set is co-sparse unless P = NP. Here is how the proof should be modified: consider the language TAUT of propositional tautologies instead of SAT. Replace property (1) by "φ is a tautology if and only if *every* formula at level ℓ is a tautology." Use $\varphi_1 \wedge \varphi_i$ instead of $\varphi_1 \vee \varphi_i$. Proceed as before. In Case 1, we can simply stop, since in this case φ_1 is not a tautology and therefore φ is not a tautology. In Case 2, we remove φ_i as before. If φ_1 is not a tautology, then removing φ_i leaves this non-tautology in the tree so we're fine. If φ_1 is a tautology, then $\varphi_i \in \text{TAUT} \Leftrightarrow \varphi_1 \wedge \varphi_i \in \text{TAUT} \Leftrightarrow \varphi_1 \wedge \varphi_j \in \text{TAUT} \Leftrightarrow \varphi_j \in \text{TAUT}$.

Remark 2.3. From the proof, we see that Mahaney's Theorem—and Fortune's Theorem, and a common generalization as in Question 3.2—has a smooth trade-off: if there is an NP-hard language with at most $s(n)$ strings of length $\leq n$, then $\text{NP} \subseteq \text{TIME}(\text{poly}(s(n)))$. In particular, if the Exponential Time Hypothesis [IP01] holds, then NP-hard sets and coNP-hard sets both have exponentially many strings of length $\leq n$, for all n.

3 Some applications

These applications are covered elsewhere, but for the sake of being self-contained and convincing the reader to actually learn and/or teach Mahaney's Theorem and Agrawal's proof thereof, I highlight them here.

3.1 Foundational questions about computational complexity

I present these first two applications in a question-and-answer format, as these are questions a bright undergraduate might ask when first learning about P versus NP. Although these questions and their answers are covered in Chapter 4 of Schöning's book [Sch86], the proofs there are slightly different. The first question is answered directly by Mahaney's Theorem; the second question is answered by a very slight variant of Agrawal's proof (quite distinct from the proof given in Schöning [Sch86, Chapter 4]).

Question 3.1. Isn't it possible that P ≠ NP, but that there is an algorithm that correctly solves SAT on all instances, but only runs in polynomial time on "most" instances, say, all but polynomially many?

This kind of question was first studied by Meyer and Paterson [MP79], who referred to an algorithm that runs in polynomial time on all instances except a sparse set as *almost polynomial-time* or *APT* (not to be confused with the more modern complexity class AlmostP, which happens to equal BPP).

Answer. No. We prove the contrapositive. Suppose there is an almost polynomial-time algorithm A that solves SAT. Let $t(n)$ be a polynomial bounding the runtime of A on all strings not in the sparse set $N \subset \Sigma^*$. Without loss of generality, we may assume that for any $x \in N$, $A(x)$ takes strictly more than $t(|x|)$ time, otherwise we could have excluded x from N in the first place. Let $N' = N \cap \text{SAT}$; we will show that if N' is empty then SAT is in P, and otherwise N' is NP-hard (in fact, NP-complete, but we won't need that).

In this latter case, since N was sparse, so is N', and hence by Mahaney's Theorem we can again conclude that $\mathsf{P} = \mathsf{NP}$.

If N' is empty, then we can directly show that SAT is in P: run $A(x)$ for $t(|x|)$ steps. If it has halted, output whatever $A(x)$ did. If it hasn't halted, reject. Hence, if N' is empty, then $\mathsf{P} = \mathsf{NP}$.

Otherwise, we show that N' is NP-hard. The following is a reduction from SAT to N': let x_{yes} be a fixed element in N'—which exists since N' is nonempty—and x_{no} a fixed element outside of N'—which exists since N' is sparse. On input x, run $A(x)$ for $t(|x|)$ steps. If $A(x)$ has accepted by this time, output x_{yes}; if $A(x)$ has rejected by this time, output x_{no}. Otherwise, simply output x. It is easily verified that this is a reduction SAT $\leq_m^p N'$. Thus, whether N' is empty or nonempty, we find that $\mathsf{P} = \mathsf{NP}$. □

Question 3.2. Isn't it possible that $\mathsf{P} \neq \mathsf{NP}$, but that there is a polynomial-time algorithm that correctly solves SAT on "most" instances, say, all but polynomially many?

A language L such that there is a polynomial-time algorithm that correctly solves L on all but polynomially many instances is called P-*close*. To rephrase this definition in terms of sparse sets: a language L is P-close if and only if there is a language $L' \in \mathsf{P}$ such that the symmetric difference $L \Delta L'$ is sparse.

Answer. No. We prove the contrapositive. Suppose that SAT is P-close: there is a language $L \in \mathsf{P}$ such that the symmetric difference $S = \text{SAT} \Delta L$ is sparse. We will show that S is NP-hard (in fact, complete) under a slightly weaker kind of reduction than many-one. A very slight variant of Agrawal's proof from above will show that this is still enough to conclude $\mathsf{P} = \mathsf{NP}$.

First we show there is a one-query polynomial-time Turing reduction from SAT to S. On input x, first compute $L(x)$ in deterministic polynomial time. If $L(x) = 0$, then $x \in \text{SAT} \Leftrightarrow x \in S$, so we make the query x to S and output the corresponding answer. If $L(x) = 1$, then $x \in \text{SAT}$ if and only if x is *not* in S, so we again make the query x to S, but now we reverse the answer.

Now we show how to modify Agrawal's proof to show that if there is a sparse NP-hard set under the preceding type of reduction (one-query polynomial-time Turing, also called "1-truth table"), then $\mathsf{P} = \mathsf{NP}$. We can think of such a reduction as follows: there is a polynomial-time function $f\colon \Sigma^* \to \Sigma^*$ which determines which query to make, and a polynomial-time function $g\colon \Sigma^* \to \{\texttt{orig}, \texttt{neg}\}$ which determines whether the answer to the query is the answer to the original question or its negation. Proceed as in Agrawal's proof, with the following modifications (we follow the notation from the proof above). Rather than pruning down to $t(n) + 1$ formulas at each level, we'll prune down to $2(t(n) + 1)$. Instead of only considering $q_i = f(\varphi_1 \vee \varphi_i)$, we consider the pair $(q_i, b_i) = (f(\varphi_1 \vee \varphi_i), g(\varphi_1 \vee \varphi_i))$.

Case 1: All q_i are distinct. Then there are either at least $t(n) + 1$ indices i such that $b_i = \texttt{orig}$ or $t(n) + 1$ indices i such that $b_i = \texttt{neg}$. In the former case, we remove φ_1 as in the original proof, with the same proof of correctness. In the latter case, we have at least $t(n) + 1$ indices i such that $b_i = \texttt{neg}$. Since there are at least $t(n) + 1$ such indices and all the q_i are distinct, at least one of the q_i with $b_i = \texttt{neg}$ is not in L. But this means exactly that $\varphi_1 \vee \varphi_i$ is satisfiable, and hence that the original formula φ is satisfiable, so we stop and accept.

Case 2: Some $q_i = q_j$ for distinct i, j. If $b_i = b_j$, then we remove φ_i as in the original proof, which works for essentially the same reason as before. If $b_i \neq b_j$, then we immediately accept. For without loss of generality, suppose $b_i = \texttt{orig}$ and $b_j = \texttt{neg}$, and let $q = q_i = q_j$. Either $q \in L$, in which case (q_i, b_i) yields a positive answer, or $q \notin L$, in which case (q_j, b_j) yields a positive answer. Either way, φ is satisfiable. □

3.2 History and original motivation: the Berman–Hartmanis Isomorphism Conjecture

When we write down an NP-complete set, we typically talk in terms of natural mathematical objects such as graphs, formulas, etc. But of course, the theory of computation demands that these objects be

encoded in some way into strings over a finite alphabet Σ. Typically we sweep this encoding process under the rug, saying that any natural and reasonable encoding will do. But technically, even amongst natural and reasonable encodings, different encodings of, say, SAT, yield *different languages* as subsets of Σ^*. However, these languages are isomorphic, in the sense that there is a polynomial-time computable bijection $f\colon \Sigma^* \to \Sigma^*$ with polynomial-time computable inverse f^{-1} which sends one encoding of SAT to another. Two subsets of Σ^* that are the same up to a polynomial-time computable and polynomial-time invertible bijection of Σ^* are called *(p-)isomorphic*. Thus, isomorphic languages are essentially just re-encodings of one another.

Berman and Hartmanis [BH77] conjectured that all NP-complete sets are isomorphic—that is, they are essentially all just re-encodings of SAT. They showed some relatively mild conditions under which an NP-complete set could be shown isomorphic to SAT, and all known natural NP-complete sets satisfy these conditions, and hence are isomorphic. (There are a few candidate exceptions [JY85, KMR95], but they were all constructed explicitly for the purpose of refuting the conjecture.) This is really quite remarkable, as it says that, for example, the set of knots of genus at most g, the set of Hamiltonian graphs, the set of solvable Super Mario Brothers levels [ADGV15], and the set of quadratic Diophantine equations with positive solutions are all not merely computationally equivalent to SAT, but simply represent different choices of encoding satisfiable formulas into Boolean strings!

The Berman–Hartmanis Isomorphism Conjecture was central to work on computational complexity in the early 1980s. One of the properties that they pointed out is preserved by isomorphism is the density of a language—the number of strings up to length n, up to polynomial transformations. It was thus natural for them to ask whether there existed an NP-complete problem that was sparse, which would therefore refute the Isomorphism Conjecture. Mahaney's Theorem was originally motivated by this question, and gave as strong as possible a negative answer. This string of ideas is also covered in the survey article by Hartmanis and Mahaney [HM80].

Sparse sets were in fact one of the central concepts in computational complexity in the 1980s, though I hope to convince you of their continued relevance and value today. In addition to the motivation from the Isomorphism Conjecture, sparse sets also became central because of their relation to circuit complexity. In particular, a language is in $\mathsf{P/poly}$ if and only if it is polynomial-time Turing reducible to some sparse set (due to Meyer, cited in Berman and Hartmanis [BH77]). From this perspective, the famous Karp–Lipton Theorem—$\mathsf{NP} \subseteq \mathsf{P/poly}$ implies $\mathsf{PH} = \Sigma_2\mathsf{P}$—can be seen as a statement about the consequences of a sparse set being NP-hard under polynomial-time Turing reductions. Improving the conclusion of Karp–Lipton to $\mathsf{P} = \mathsf{NP}$—or equivalently, improving the reductions in Mahaney's Theorem to polynomial-time Turing reductions—requires nonrelativizing techniques [IM89] (the same is even true for strengthening Karp–Lipton to conclude $\mathsf{PH} = \mathsf{P}^{\mathsf{NP}}$ [Hel86]). However, Mahaney's Theorem *has* been extended from polynomial-time many-one reductions to "bounded truth-table" reductions: nonadapative Turing reductions that make only $O(1)$ queries [OW91], and even slightly beyond [AKM92]. For more on the history and surveys of related results, see, e. g., [HOW92, You92a, You92b, GH00a, GH00b, CO97] and references therein.

3.3 Representations of the symmetric group and geometric complexity theory

We conclude with a much more recent application [IMW15], which is a statement about the representation theory of the symmetric groups, which was, in fact, not known prior to this application of a sparsity theorem. Although fully understanding the technical details here would require reading some other papers and having some knowledge of representation theory, the value of Mahaney's Theorem for representation theory and geometric complexity theory (GCT) can be understood without such knowledge. We try to give an exposition of this application to make such appreciation possible, without knowing the relevant representation theory. (The relevant representation theory can be found in, e. g., the first part of Fulton and Harris [FH91], Fulton's book on Young tableaux [Ful97], or James's book [Jam78].)

The GCT Program (see, e. g., [MS01, MS08, Mul11a, Mul11b, Mul12] and references therein and thereto) proposes to resolve major conjectures like Permanent versus Determinant (VP_{ws} versus VNP) or P versus NP by associating to each complexity class certain integer multiplicities coming from representation theory. The connection they made to lower bounds was the following:

Theorem 3.3 (Mulmuley and Sohoni [MS01])**.** *If there is some representation-theoretic multiplicity which is larger for the permanent than for the determinant, then* $\mathsf{VP}_{ws} \neq \mathsf{VNP}$.

They also made the following conjecture, which is formally stronger than what is needed to use the preceding result:

Conjecture 3.4 (Mulmuley and Sohoni [MS01])**.** *There exist representation-theoretic multiplicities that are* zero *for* VP_{ws} *and nonzero for* VNP.

Although this conjecture was recently disproven [BIP16], the following result is still a nice application of computational complexity to representation theory.[1] The so-called "Kronecker coefficients" are some of the relevant representation-theoretic multiplicities associated to VP_{ws} or the determinant (they are, more precisely, an upper bound on the multiplicities associated to the determinant). The Kronecker coefficients are quite classical mathematical objects, going back more than 100 years. **For those who know a little representation theory:** more precisely, the Kronecker coefficients are the tensor product multiplicities for the symmetric group S_n. In other words, given two irreducible representations V_λ, V_μ of S_n, specified by partitions λ, μ of n, the Kronecker coefficient $k_{\lambda,\mu,\nu}$ is the multiplicity of the irreducible representation V_ν in the decomposition of $V_\lambda \otimes V_\mu$ into a direct sum of irreducible representations.

Theorem 3.5 (Ikenmeyer, Mulmuley, Walter [IMW15])**.** *If* $\mathsf{P} \neq \mathsf{NP}$*, then the Kronecker coefficients are zero super-polynomially often, and hence, so are the multiplicities associated to the determinant and* VP_{ws}.

In fact, they show the conclusion of this theorem unconditionally, but they do so by showing that deciding the positivity of Kronecker coefficients is NP-hard, and then reducing from an NP-complete problem that is known to have exponentially many no-instances [IMW15, Lemma 5.1 and Theorem 5.2].

Proof. Bürgisser and Ikenmeyer [BI13] showed that there is an NP-complete problem L, whose naturally associated #P-complete counting problem is an upper bound on the Kronecker coefficients. In particular, the number of no-instances of L at length n is a lower bound on the number of Kronecker coefficients of length n that are zero. By Fortune's Theorem (see the simple proof in Remark 2.2), if $\mathsf{P} \neq \mathsf{NP}$, then the number of no-instances of L must grow super-polynomially in n. □

Remark 3.6. Because of the smooth tradeoff in Fortune's Theorem (see Remark 2.3), even without the unconditional result of [IMW15], from their conditional result one immediately gets that if the Exponential Time Hypothesis [IP01] holds, then the Kronecker coefficients are zero exponentially often (which is what they then proved unconditionally).

[1] For historical purposes, we note that before the recent disproof of the conjecture, analogous and related results (see, e. g., Ikenmeyer [Ike12, Section 6.3]) as well as computer experiments (also by Ikenmeyer, summarized succinctly in, e. g., Landsberg [Lan15, Section 6.6]) suggested that only very few multiplicities associated to VP_{ws} were zero, which would have made it very difficult to prove the above conjecture. Of course, it's expected to be difficult given the nature of the questions it is approaching, but if too many of these multiplicities were zero, it might have made them *very* hard to find. This result suggested that, despite the numerical evidence, in fact lots of these multiplicities were zero. The recent disproof of the conjecture shows that the associated multiplicities for VNP are also zero.

References

[AA96] Manindra Agrawal and Vikraman Arvind. Geometric sets of low information content. *Theoret. Comput. Sci.*, 158(1-2):193–219, 1996. doi:10.1016/0304-3975(95)00073-9.

[ADGV15] Greg Aloupis, Erik D. Demaine, Alan Guo, and Giovanni Viglietta. Classic Nintendo games are (computationally) hard. *Theoret. Comput. Sci.*, 586:135–160, 2015. Originally appeared in *Fun With Algorithms*, Springer Lec. Notes Comp. Sci., 8496:40–51; preprint available as arXiv:1203.1895 [cs.CC]. doi:10.1016/j.tcs.2015.02.037.

[AKM92] Vikraman Arvind, Johannes Köbler, and Martin Mundhenk. On bounded truth-table, conjunctive, and randomized reductions to sparse sets. In *Foundations of software technology and theoretical computer science (New Delhi, 1992)*, volume 652 of *Lecture Notes in Computer Science*, pages 140–151. Springer, Berlin, 1992. doi:10.1007/3-540-56287-7_101.

[BH77] Leonard Berman and Juris Hartmanis. On isomorphisms and density of NP and other complete sets. *SIAM J. Comput*, 6(2):305–322, 1977. Originally appeared in STOC '76. doi:10.1137/0206023.

[BI13] Peter Bürgisser and Christian Ikenmeyer. Explicit lower bounds via geometric complexity theory. In *STOC '13: 45th Annual ACM Symposium on Theory of Computing*, 2013. Also available as arXiv:1210.8368 [cs.CC]. doi:10.1145/2488608.2488627.

[BIP16] Peter Bürgisser, Christian Ikenmeyer, and Greta Panova. No occurrence obstructions in geometric complexity theory. In *FOCS '16: 57th Annual IEEE Symposium on Foundations of Computer Science*, 2016. Also available as arXiv:1604.06431 [cs.CC].

[CO97] Jin-Yi Cai and Mitsunori Ogihara. Sparse sets versus complexity classes. In *Complexity theory retrospective, II*, pages 53–80. Springer, New York, 1997.

[FH91] William Fulton and Joe Harris. *Representation theory*, volume 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991. doi:10.1007/978-1-4612-0979-9.

[For79] Steven Fortune. A note on sparse complete sets. *SIAM J. Comput*, 8(3):431–433, 1979. doi:10.1137/0208034.

[Ful97] William Fulton. *Young tableaux*, volume 35 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1997. With applications to representation theory and geometry.

[GH00a] Christian Glaßer and Lane A. Hemaspaandra. A moment of perfect clarity I: The parallel census technique. *SIGACT News*, 31(3):37–42, September 2000. doi:10.1145/356458.356459.

[GH00b] Christian Glaßer and Lane A. Hemaspaandra. A moment of perfect clarity II: Consequences of sparse sets hard for NP with respect to weak reductions. *SIGACT News*, 31(4):39–51, December 2000. doi:10.1145/369836.369838.

[Hel86] Hans Heller. On relativized exponential and probabilistic complexity classes. *Inform. and Control*, 71(3):231–243, 1986. doi:10.1016/S0019-9958(86)80012-2.

[HM80] Juris Hartmanis and Stephen R. Mahaney. An essay about research on sparse NP complete sets. In *MFCS '80: 9th Symposium on Mathematical Foundations of Computer Science*, volume 88 of *Lecture Notes in Computer Science*, pages 40–57. Springer, Berlin-New York, 1980. Preprint available as Cornell Comp. Sci. Tech. Report TR80-422. doi:10.1007/BFb0022494.

[HOW92] Lane A. Hemachandra, Mitsunori Ogiwara, and Osamu Watanabe. How hard are sparse sets? In *Proceedings of the Seventh Annual Structure in Complexity Theory Conference (Boston, MA, 1992)*, pages 222–238. IEEE Comput. Soc. Press, Los Alamitos, CA, 1992. doi:10.1109/SCT.1992.215396.

[Ike12] Christian Ikenmeyer. *Geometric complexity theory, tensor rank, and Littlewood–Richardson coefficients*. PhD thesis, Institute of Mathematics, University of Paderborn, 2012. URL: http://math-www.uni-paderborn.de/agpb/work/ikenmeyer_thesis.pdf.

[IM89] Neil Immerman and Stephen R. Mahaney. Relativizing relativized computations. *Theoret. Comput. Sci.*, 68(3):267–276, 1989. doi:10.1016/0304-3975(89)90164-3.

[IMW15] Christian Ikenmeyer, Ketan D. Mulmuley, and Michael Walter. On vanishing of Kronecker coefficients. arXiv:1507.02955 [cs.CC], 2015.

[IP01] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. System Sci.*, 62(2):367–375, 2001. Special issue on the Fourteenth Annual IEEE Conference on Computational Complexity (Atlanta, GA, 1999). doi:10.1006/jcss.2000.1727.

[Jam78] G. D. James. *The representation theory of the symmetric groups*, volume 682 of *Lecture Notes in Mathematics*. Springer, Berlin, 1978.

[JY85] Deborah Joseph and Paul Young. Some remarks on witness functions for nonpolynomial and noncomplete sets in NP. *Theoret. Comput. Sci.*, 39(2-3):225–237, 1985. doi:10.1016/0304-3975(85)90140-9.

[KMR95] Stuart A. Kurtz, Stephen R. Mahaney, and James S. Royer. The isomorphism conjecture fails relative to a random oracle. *J. Assoc. Comput. Mach.*, 42(2):401–420, 1995. doi:10.1145/201019.201030.

[Lan15] J. M. Landsberg. Geometric complexity theory: an introduction for geometers. *Ann. Univ. Ferrara Sez. VII Sci. Mat.*, 61(1):65–117, 2015. Preprint available as arXiv:1305.7387 [math.AG]. doi:10.1007/s11565-014-0202-7.

[Mah82] Stephen R. Mahaney. Sparse complete sets for NP: solution of a conjecture of Berman and Hartmanis. *J. Comput. System Sci.*, 25(2):130–143, 1982. doi:10.1016/0022-0000(82)90002-2.

[MP79] Albert R. Meyer and Michael S. Paterson. With what frequency are apparently intractable problems difficult? Technical Report MIT/LCS/TM-126, Massachusetts Institute of Technology, Laboratory for Computer Science (LCS), 1979. URL: http://www.dcs.warwick.ac.uk/people/academic/Mike.Paterson/papers/MIT-LCS-TM-126.pdf.

[MS01] Ketan D. Mulmuley and Milind Sohoni. Geometric complexity theory I: an approach to the P vs. NP and related problems. *SIAM J. Comput*, 31(2):496–526, 2001. doi:10.1137/S009753970038715X.

[MS08] Ketan D. Mulmuley and Milind Sohoni. Geometric complexity theory. II. Towards explicit obstructions for embeddings among class varieties. *SIAM J. Comput*, 38(3):1175–1206, 2008. doi:10.1137/080718115.

[Mul11a] Ketan D. Mulmuley. Geometric Complexity Theory VI: the flip via positivity. Technical report, University of Chicago, 2011. URL: http://gct.cs.uchicago.edu/gct6.pdf.

[Mul11b] Ketan D. Mulmuley. On P vs. NP and Geometric Complexity Theory. *J. Assoc. Comput. Mach.*, 58(2):Art. 5, 2011. doi:10.1145/1944345.1944346.

[Mul12] Ketan D. Mulmuley. The GCT program toward the P vs. NP problem. *Commun. ACM*, 55(6):98–107, June 2012. doi:10.1145/2184319.2184341.

[OW91] Mitsunori Ogiwara and Osamu Watanabe. On polynomial-time bounded truth-table reducibility of NP sets to sparse sets. *SIAM J. Comput*, 20(3):471–483, 1991. doi:10.1137/0220030.

[Sch86] Uwe Schöning. *Complexity and structure*, volume 211 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1986. doi:10.1007/3-540-16079-5.

[You92a] Paul Young. How reductions to sparse sets collapse the polynomial-time hierarchy: A primer; part I: Polynomial-time Turing reductions. *SIGACT News*, 23(3):107–117, June 1992. doi:10.1145/141914.141921.

[You92b] Paul Young. How reductions to sparse sets collapse the polynomial-time hierarchy: A primer: Part II restricted polynomial-time reductions. *SIGACT News*, 23(4):83–94, October 1992. doi:10.1145/148080.148084.